

NORMALIZING TOPOLOGICALLY MINIMAL SURFACES III: BOUNDED COMBINATORICS

DAVID BACHMAN

ABSTRACT. We show that there are a finite number of possible pictures for a surface in a tetrahedron with local index n . Combined with previous results, this establishes that any topologically minimal surface can be transformed into one with a particular normal form with respect to any triangulation.

1. INTRODUCTION

A surface in a triangulated 3-manifold is *normal* if it intersects each tetrahedron transversally in a collection of flat triangles and quadrilaterals. H. Kneser introduced normal surfaces to show that the prime decomposition of a 3-manifold is finite [Kne29]. Many years later, W. Haken used normal surfaces to give the first algorithm to determine if a knot is the unknot, and to show that any collection of pairwise disjoint, non-parallel incompressible surfaces is finite [Hak61] (see also [Bac04]). The key to this latter proof was to show that any incompressible surface can be isotoped to a normal surface.

In 1993, Rubinstein introduced the idea of an *almost normal* surface. Such a surface meets each tetrahedron of a triangulated 3-manifold in a collection of normal disks (triangles or quadrilaterals), with the single exception of a disk whose boundary meets the 1-skeleton eight times (i.e. an octagon) or two normal disks connected by a single unknotted tube. Rubinstein [Rub95a] and Thompson [Tho94] used almost normal surfaces to produce an algorithm to determine if a given 3-manifold is S^3 . Subsequently Rubinstein [Rub95b] and Stocking [Sto00] used almost normal surfaces to produce an algorithm to determine the Heegaard genus of a given 3-manifold. More recently, Li used almost normal surfaces to prove the generalized Waldhausen conjecture [Li07]: any closed, atoroidal 3-manifold admits finitely many Heegaard splittings of each genus. The key to these results is to show that any *strongly*

Date: March 28, 2013.

Partially supported by NSF grant DMS-1207804.

irreducible [CG87] Heegaard splitting is isotopic to an almost normal surface.

These results illustrate the idea that a “normal form” for a particular class of surfaces with respect to an arbitrary triangulation enables proofs of both algorithmic and finiteness results. In [Bac10], the author introduced the idea of *topologically minimal surfaces* as a way of generalizing both incompressible and strongly irreducible surfaces. Such surfaces have an *index*, with index 0 surfaces being incompressible and index 1 surfaces being strongly irreducible. Given the aforementioned results, the obvious question is, “Can every topologically minimal surface be isotoped to some normal form?” Resolving this question is the purpose of the present series of papers.

In the first paper in this series [Baca] we reduced the problem to a local one. In that work, we showed that a topologically minimal surface with index n in a triangulated 3-manifold can be transformed¹ to one that meets each tetrahedron in a surface with a well-defined *local index*. Furthermore, the sum of these local indices, taken over all tetrahedra, will be at most n . What remains, then, is to classify all surfaces in a tetrahedron with well-defined local index.

In the second paper in this series [Bacb] we obtained a complete classification of simply connected surfaces in a tetrahedron with local index n . There we showed that such a surface must look like a helicoid, and the boundary of such a surface will meet the 1-skeleton precisely $4(n + 1)$ times. This result bears a striking resemblance with Colding and Minicozzi’s classification of the local picture of a geometrically minimal surface that is simply connected [CM04a], [CM04b], [CM04c], [CM04d].

In this final paper, we study the general case of surfaces in a tetrahedron Δ with local index n . Our goal here is to show that there is a finite number of possibilities for such a surface. The results of this paper can all be summarized in the following theorem.

Theorem 1.1. *Let H be a surface in a tetrahedron Δ with local index n . Then*

- (1) *H is unknotted and the components of H are not nested.*
- (2) *$\sum 1 - \chi(H_i) \leq n$, where the sum is taken over all components H_i of H .*
- (3) *Each component of H has well-defined local index, and the sum of these indices is precisely n .*

¹Such a transformation may involve isotopy, ∂ -compression, and some compressions. Compressions that arise are very constrained, only happen when n is at least three, and reduce the index.

- (4) If H' is a component of H then $|\partial H'| \leq 4(n' + 1)$, where n' is the local index of H' .

The results of this theorem imply that there are a finite number of possibilities for a surface in a tetrahedron with local index n . For example, when $n = 0$ the theorem implies that the surface H must be a collection of disks with boundary length 3 or 4. This recovers the proof that any incompressible surface can be isotoped to a normal surface. When $n = 1$ the above result implies that there is at most one component of H that is not a normal triangle or quadrilateral. This component can either be a disk or an annulus, and the total length of its boundary is at most 8. Hence, the exceptional component is either a disk with boundary length 8, or an unknotted annulus connecting two loops that have length 3 or 4. Taken with the results from the prequels, this recovers Rubinstein and Stocking's result that any strongly irreducible surface can be isotoped to an almost normal one.

The index 2 case is of particular interest, as it has immediate consequences that were not previously known. Theorem 1.1 implies that a (connected) surface in a tetrahedron with local index 2 is either:

- (1) A disk with boundary of length 12 (i.e. a 12-gon).
- (2) Three normal disks, connected by two unknotted tubes, to form a pair-of-pants.
- (3) An octagon connected to a normal triangle by an unknotted tube.
- (4) An octagon connected to itself by an unknotted tube to form a punctured torus.
- (5) A 12-gon connected to itself by an unknotted tube to form a punctured torus.

Although Theorem 1.1 does not rule out this last surface as a possible index 2 picture, we conjecture that it is not. This illustrates an important point: Theorem 1.1 only gives a finite set of surfaces that contains all possible ones that have local index n . However, it is likely the case that not all of the surfaces in this set will have the desired index. Despite this, the theorem does establish that for each local index the possibilities are finite, which is what is necessary for applications. For example taken with Theorem 1.2 of [Baca] and Lemma 2.10 of [BDTS], we obtain the following result:

Corollary 1.2. *Let M be a compact, irreducible 3-manifold with incompressible, connected, toroidal boundary. Then there is a finite set of slopes Ω on ∂M such that the boundary slope of any non-peripheral, topologically minimal, index n surface in M is at a distance of at most n (as measured in the Farey graph) from Ω .*

Proof. Choose any triangulation \mathcal{T} of M . Let H be a topologically minimal surface of index n in M . By Theorem 1.2 of [Baca], if H is not peripheral then we can isotope, compress and ∂ -compress H to a surface H' such that the local indices of $H' \cap \Delta_i$ sum to at most n , where $\{\Delta_i\}$ are the tetrahedra of \mathcal{T} . Furthermore, the slope of $\partial H'$ will be at most distance n (as measured in the Farey graph) from the slope of ∂H (see Remark 2.18 of [Baca]). By Theorem 1.1, for each i , each component of the surface $H' \cap \partial \Delta_i$ is of bounded length, and thus H' is in one of a finite number of *compatibility classes* (see Definition 2.5 of [BDTS]). Lemma 2.10 in [BDTS] implies that each such compatibility class determines a unique slope on ∂M (see in particular the paragraph following the proof of this lemma). \square

This result, together with Theorem 4.9 of [Bac10] then implies the following, which says roughly that if manifolds X and Y with toroidal boundary are glued together by a “sufficiently complicated” map, then no new index n surfaces are created.

Corollary 1.3. *Let X and Y be compact, irreducible 3-manifolds with incompressible, connected, toroidal boundary. Let $\phi : \partial X \rightarrow \partial Y$ be an homeomorphism, and $\psi : \partial Y \rightarrow \partial Y$ be any Anosov map. Then there is a value m_0 depending only on n such that for all $m \geq m_0$, any topologically minimal, index n surface in the manifold $X \cup_{\psi^m \circ \phi} Y$ is isotopic into X or Y .*

Proof. Let Ω_X and Ω_Y be the finite sets of slopes on ∂X and ∂Y guaranteed by Corollary 1.2, so that if $H_X \subset X$ and $H_Y \subset Y$ are surfaces with topological index n , then their boundaries are at most distance n from Ω_X and Ω_Y , respectively. Choose m_0 so that if $m \geq m_0$, then the set Ω_Y is at least distance $2n + 1$ from the set $\psi^m \circ \phi(\Omega_X)$ in the Farey graph of ∂Y .

Now suppose H is an index n surface in $M = X \cup_{\psi^m \circ \phi} Y$ that is not isotopic into either X or Y . Let T be the image of ∂Y in this manifold. By Theorem 4.9 of [Bac10], H may be isotoped to meet T transversally away from p saddle points, so that the sum of the indices of the components of $H - N(T)$ is at most $n - p$. In particular, there are components $H_X \subset X$ and $H_Y \subset Y$ of $H - N(T)$ whose indices are at most n . Furthermore, ∂H_X and ∂H_Y are essential on T , and come from $H \cap T$ by resolving p saddle points. It follows that these loops are disjoint, and since T is a torus, must actually be the same.

We now apply Corollary 1.2, which says that ∂H_X is at a distance of at most n from Ω_X , and ∂H_Y is distance at most n from Ω_Y . Since

$\psi^m \circ \phi(\partial H_X) = \partial H_Y$, we conclude the distance between Ω_Y and $\psi^m \circ \phi(\Omega_X)$ is at most $2n$, contradicting our choice of m_0 . \square

In Corollary 11.2 of [Bac13] we constructed the first example of a manifold that contains a non-minimal genus, unstabilized, weakly reducible Heegaard splitting. By using Corollary 1.3, we can improve this construction to show that there is a manifold that contains an infinite number of such splittings.

Corollary 1.4. *There is a closed 3-manifold M that contains infinitely many non-isotopic, unstabilized Heegaard splittings that are weakly reducible (i.e. splittings whose Hempel distance is precisely one).*

Proof. In the parlance of [Bac13], Corollary 1.3 establishes that if X and Y are manifolds with toroidal boundary that are glued by a sufficiently complicated homeomorphism, then the torus T at their interface is a “ g -barrier surface” for all index 0, 1, and 2 surfaces, regardless of the genus g . Thus, the conclusion of Theorem 11.1 of [Bac13] holds, regardless of the genus of the surface H_i . The construction given in Corollary 11.2 of [Bac13] then works for every genus. \square

Acknowledgements. The author is immensely grateful to Ryan Derby-Talbot and Eric Sedgwick for invaluable conversations during the preparation of this work.

2. TOPOLOGICALLY MINIMAL SURFACES IN TETRAHEDRA

Throughout this paper H will represent a properly embedded surface in a tetrahedron Δ . The 1- and 2-skeleton of Δ will be denoted \mathcal{T}^1 and \mathcal{T}^2 , respectively.

Definition 2.1. A compressing disk for H is a disk C such that $C \cap H$ is an essential loop on H . We will use the notation $\mathcal{C}(H)$ to refer to the set of compressing disks for H .

Definition 2.2. An edge-compressing disk for H is a disk E such that $\partial E = \alpha \cup \beta$, where $E \cap H = \alpha$ and $E \cap \mathcal{T}^1 = \beta$ is a subarc of an edge of \mathcal{T}^1 . We will use the notation $\mathcal{E}(H)$ to refer to the set of edge-compressing disks for H .

Definition 2.3. We let $\mathcal{CE}(H) = \mathcal{C}(H) \cup \mathcal{E}(H)$.

Definition 2.4. For any disk $D \in \mathcal{CE}(H)$, we define H/D to be the surface obtained from H as follows. Let $N(D)$ denote a submanifold of Δ homeomorphic to $D \times I$, so that $D = D \times \{\frac{1}{2}\}$, $N(D) \cap H = (D \cap H) \times I$, and $N(D) \cap \partial\Delta = (D \cap \partial\Delta) \times I$. Then H/D is obtained from H by removing $N(D) \cap H$ and replacing it with $D \times \partial I$. The

two subdisks $D \times \partial I$ in H/D will be called the *scars* of D . When \mathcal{D} is a collection of pairwise disjoint disks in $\mathcal{CE}(H)$, then H/\mathcal{D} will denote the result of simultaneously performing the above operation on each individual disk in \mathcal{D} .

We define an equivalence relation as follows:

Definition 2.5. Disks $D, D' \in \mathcal{CE}(H)$ are *equivalent*, $D \sim D'$, if D and D' are isotopic through disks in $\mathcal{CE}(H)$.

Definition 2.6. The *disk complex* $[\mathcal{CE}(H)]$ is the complex defined as follows: vertices correspond to equivalence classes $[D]$ of the disk set $\mathcal{CE}(H)$. A collection of n vertices spans an $(n - 1)$ -simplex if there are representatives of each which are pairwise disjoint away from a neighborhood of \mathcal{T}^1 .

Definition 2.7. The surface H is said to be *topologically minimal* in the tetrahedron Δ if the complex $[\mathcal{CE}(H)]$ is either empty or non-contractible. In the former case we define the *local index* of H to be 0. In the latter case, the *local index* of H is defined to be the smallest n such that $\pi_{n-1}([\mathcal{CE}(H)])$ is non-trivial.

The following definition and lemma will aid us in establishing the third conclusion of Theorem 1.1.

Definition 2.8. We say two components of a surface $H \subset \Delta$ are *normally separated* if there is a normal triangle or quadrilateral in Δ separating them. The *normal components* of H are then the disjoint unions of the components of H that are not normally separated. A surface is *normally connected* if it has one normal component. (Equivalently, no two components of H can be normally separated.)

For example, if the components of H are two parallel octagons and two normal triangles, then H has three normal components. One of these normal components will contain the two octagons, and the other two will each contain a single normal triangle.

Lemma 2.9. *If a surface H in a tetrahedron Δ has local index n , then the normal components of H have well-defined local indices, and the sum of these indices is precisely n .*

Proof. Suppose $H = H_1 \cup H_2$, where H_1 and H_2 are subsurfaces that are separated by a normal triangle or quadrilateral. Then every compressing and edge-compressing disk for H is a compressing or edge compressing disk for either H_1 or H_2 , and vice versa. Hence, the complex $[\mathcal{CE}(H)]$ is the join of the complexes $[\mathcal{CE}(H_1)]$ and $[\mathcal{CE}(H_2)]$. It follows that both H_1 and H_2 have well-defined local indices, and the

sum of these indices is precisely the local index of H (see, for example, the proof of Theorem 4.7 of [Bac10]). The lemma now follows by inducting on the number of normal components of H . \square

Following this, to establish the third conclusion of Theorem 1.1 we only need to show that a normally connected surface H with well-defined local index is actually connected. Hence, in the remaining sections we will assume the surface H under consideration is normally connected. In particular, H has no components that are normal triangles or quadrilaterals.

3. THE TOPOLOGY OF TOPOLOGICALLY MINIMAL SURFACES IN TETRAHEDRA

In this section we limit the topology of a topologically minimal surface H in a tetrahedron Δ and establish the first two conclusions of Theorem 1.1.

Lemma 3.1. *Let H be a surface in a tetrahedron Δ . If \mathcal{E} is a collection of pairwise disjoint edge-compressing disks for H , and C is a compressing disk for H/\mathcal{E} , then C is a compressing disk for H that is disjoint from every element of \mathcal{E} .*

Proof. To reconstruct H from H/\mathcal{E} , we attach bands between subarcs of $\partial H/\mathcal{E}$. As C is in the interior of Δ , and these bands can be chosen to be arbitrarily close to $\partial\Delta$, the result follows. \square

The first two conclusions of Theorem 1.1 follow immediately from the following theorem.

Theorem 3.2. *Suppose H is an index n , topologically minimal surface in a tetrahedron Δ . Then there is a collection \mathcal{E} of at most n edge-compressing disks for H that are pairwise disjoint away from a neighborhood of \mathcal{T}^1 , such that H/\mathcal{E} is a collection of disks.*

Proof. We will say a loop of ∂H is *non-trivial* if it does not bound a disk component of H . If H is a collection of disks, the result is trivial. If not, then a component of ∂H that is innermost on $\partial\Delta$ among all non-trivial loops bounds a compressing disk for H . Thus, $\mathcal{CE}(H) \neq \emptyset$. By definition then, $n \geq 1$ and there is a homotopically non-trivial map ι of an $(n-1)$ -sphere S into $[\mathcal{CE}(H)]$. Triangulate S so that the map ι is simplicial.

Claim 3.3. *There is a vertex v of S such that $\iota(v)$ represents an edge-compressing disk.*

Proof. If not, then every vertex in the image of ι represents a compressing disk. As above, a component of ∂H that is innermost on $\partial\Delta$ among all non-trivial loops bounds a compressing disk C for H . Furthermore, the disk C can be isotoped to be disjoint from every other compressing disk for H . Thus, the vertex $[C]$ of $[\mathcal{CE}(H)]$ is connected by an edge to every vertex in the image of ι . This contradicts the assumption that ι is homotopically non-trivial. \square

We will call a simplex of S an *edge-simplex* if it is spanned by vertices v such that $\iota(v)$ is an edge-compressing disk. Henceforth, we will assume that all choices have been made so that the total number ϵ of edge-simplices of S (of any dimension) is as small as possible. It follows from the previous claim that $\epsilon > 0$.

Claim 3.4. *Suppose σ is a maximal dimensional edge-simplex of S (i.e. σ is not a face of any other edge-simplex). Let \mathcal{E} be a collection of disks representing the vertices of $\iota(\sigma)$. Then H/\mathcal{E} is incompressible.*

Proof. By way of contradiction, assume H/\mathcal{E} is compressible. Let α be a component of $\partial(H/\mathcal{E})$ that is innermost on $\partial\Delta$, among all loops that do not bound disk components of H/\mathcal{E} . Then α bounds a compressing disk C for H/\mathcal{E} that can be made disjoint from every other compressing disk for H/\mathcal{E} . By Lemma 3.1, C is also a compressing disk for H that is disjoint from every disk representing the vertices of $\iota(\sigma)$.

Let Σ denote the union of the simplices in S that contain σ . Thus, each vertex v of a simplex of Σ is either contained in σ , or is connected by an edge to every vertex of σ . In the latter case, $\iota(v)$ must be represented by a compressing disk V for H , as otherwise $\sigma \cup v$ would be an edge-simplex containing σ , contradicting our assumption that σ was of maximal dimension. As v is connected by an edge to each vertex of σ it must be disjoint from every edge-compressing disk in \mathcal{E} . Thus, the disk V is also a compressing disk for the surface H/\mathcal{E} . Since C is disjoint from every compressing disk for H/\mathcal{E} , it must be the case that either $[C] = [V]$ in $[\mathcal{CE}(H)]$, or $[C]$ is connected to $[V]$ by an edge. In either case, this shows that we can alter the map ι by replacing Σ with the cone of $\partial\Sigma$ to a vertex c , where $\iota(c) = [C]$. This replacement removes the edge-simplex σ without introducing any new edge-simplices, thus lowering ϵ . \square

The theorem now follows from the previous claims. \square

The following lemmas are an important step in the proof of the third conclusion of Theorem 1.1.

Lemma 3.5. *Suppose H is a normally connected surface in a tetrahedron Δ with well-defined local index and A is a component of $\partial\Delta - \partial H$. If there is a loop α in A of length three or four, then there are boundary components of A on either side of α that belong to the same component of H .*

Proof. Let C be a normal triangle or quadrilateral bounded by α . Isotope H so that $|C \cap H|$ is minimal. Let \mathcal{E} denote a collection of edge-compressing disks for H such that H/\mathcal{E} is a collection of disks, as guaranteed to exist by Theorem 3.2. We claim that some element $E \in \mathcal{E}$ meets $\alpha = \partial C$. If not, then there will be a non-trivial loop in $C - (H \cup \mathcal{E})$. Let β denote an innermost loop of a maximal collection of pairwise disjoint, non-parallel, non-trivial loops in $C - (H \cup \mathcal{E})$. As H/\mathcal{E} is a collection of disks, the loop β must bound a disk D in $\Delta - H/\mathcal{E}$. The surface H can be recovered from H/\mathcal{E} by attaching bands that can be chosen to be arbitrarily close to $\partial\Delta$. In particular, these bands can be chosen to be disjoint from D . We conclude the disk D is disjoint from H . It follows that isotoping D into C reduces $|H \cap C|$, contradicting minimality. We conclude some element $E \in \mathcal{E}$ meets α . The arc $E \cap H$ will then be a path in H connecting points of ∂A on opposite sides of α . \square

Lemma 3.6. *If H is a normally connected surface in a tetrahedron Δ with well-defined local index and every component of ∂H has length three or four, then H is connected.*

Proof. Each component A of $\partial\Delta - \partial H$ is either an annulus, “pair-of-pants” (i.e. thrice-punctured sphere), or “t-shirt” (i.e. four times punctured sphere). It suffices to show that in all cases, every component of ∂A is on the same component of H .

Consider first the case when A is an annulus. Then a loop α on A that is normally parallel to either boundary component will have length three or four. Applying Lemma 3.5 to the loop α then tells us that both components of ∂A belong to the same component of H .

When A is a pair-of-pants, let A_1 , A_2 , and A_3 denote the boundary components of A , and α_i a loop in A that is normally parallel to A_i . Applying Lemma 3.5 to α_1 , we conclude that A_1 is on the same component of H as either A_2 or A_3 . Suppose the former. We then apply Lemma 3.5 to the loop α_3 to conclude A_3 is on the same component of H as A_1 and A_2 . Thus, all three boundary components of A are on the same component of H .

We now deal with this case when A is a “t-shirt”. This arises when all three components of ∂A are length three. Let A_1 , A_2 , A_3 , and A_4 denote the boundary components of A , and α_i a loop in A that is

normally parallel to A_i . Applying Lemma 3.5 to α_1 , we conclude that A_1 is on the same component of H as either A_2 , A_3 , or A_4 . Without loss of generality, assume A_2 . We now apply Lemma 3.5 to the loop α_3 to conclude A_3 is on the same component of H as A_1 , A_2 , or A_4 . There are two cases to consider:

Case 1. A_3 is on the same component of H as A_1 and A_2 . Then apply Lemma 3.5 to the loop α_4 to conclude A_4 is on the same component of H as A_1 , A_2 , and A_3 .

Case 2. A_3 is on the same component of H as A_4 . Let α be a loop on A which separates A_1 and A_2 from A_3 and A_4 . Then α has length four, so we may apply Lemma 3.5 to it. We conclude A_1 and A_2 are on the same component of H as A_3 and A_4 . \square

4. BOUNDING THE LENGTH OF BOUNDARY CURVES WITH NO TRIANGULAR COMPONENTS

In this section we establish the fourth conclusion of Theorem 1.1, assuming ∂H contains only loops of length at least 8. In this case, it is well known that \mathcal{T}^1 , the 1-skeleton of Δ , consists of three pairs of opposite edges, s and s' , m and m' , and l and l' , such that


- (1) $|\partial H \cap s| = |\partial H \cap s'| = a > 0$
- (2) $|\partial H \cap m| = |\partial H \cap m'| = b \geq a$
- (3) $|\partial H \cap l| = |\partial H \cap l'| = c = a + b$

The length of ∂H is precisely $2(a + b + c) = 4c$. Hence, an upper bound on c provides us with an upper bound on $|\partial H|$. We will call s and s' the *short edges*, m and m' the *medium edges*, and l and l' the *long edges*. Note that when $a = b$ the choice of short and medium edges can be made arbitrarily.

For the remainder of this section v will denote the vertex of Δ where the edges s , m , and l meet. For any edge e incident to the vertex v , we can identify the frontier R_e of a neighborhood $N(e)$ of e in Δ with $I \times I$ so that

- for each $t \in I$, the arc $t \times I$ is parallel in $N(e)$ to e , and
- for each component γ of $H \cap R$ there is an $s \in I$ such that $\gamma = I \times s$.

Definition 4.1. Let τ_{left}^e and τ_{right}^e denote the 2-simplices of Δ that are incident to the edges $0 \times I$ and $1 \times I$ of R_e , respectively. We say these are the 2-simplices that are to the *left* and *right* of e . For the three edges s , m , and l , we assume identifications of R_s , R_m , and R_l


 FIGURE 1. Defining left and right *switches*.

with $I \times I$ have been chosen so that $\tau_{\text{right}}^m = \tau_{\text{left}}^l$, $\tau_{\text{right}}^l = \tau_{\text{left}}^s$, and $\tau_{\text{right}}^s = \tau_{\text{left}}^m$.

Definition 4.2. A subregion of R_e cobounded by arcs of $H \cap R_e$, which does *not* meet a region of τ_{left}^e bounded by parallel normal arcs of ∂H is called a *left switch* of e . See Figure 1. Right switches are defined similarly.

Definition 4.3. Let E denote an edge-compressing disk for H incident to an edge e . Suppose the arcs of $H \cap (R_e \cup \tau_{\text{left}}^e)$ that meet E are parallel, as in Figure 2. Let E' denote the union of the shaded region in the figure, together with the disk E . Then E' is an edge-compressing disk for H , called the *push-off* of E to its left. Push-offs to the right are defined similarly.

The push-off operation can sometimes be iterated, as follows:

Definition 4.4. Let E denote an edge-compressing disk for H incident to an edge e . Let f denote an edge of Δ , so that $\tau_{\text{right}}^f = \tau_{\text{left}}^e$. Suppose the arcs of $H \cap (R_e \cup \tau_{\text{left}}^e \cup R_f \cup \tau_{\text{left}}^f)$ that meet E are parallel, as in Figure 3. Let E'' denote the union of the shaded region in the figure, together with the disk E . Then E'' is an edge-compressing disk for H , called the *second push-off* of E to its left. Second push-offs to the right are defined similarly.

Definition 4.5. Let \mathcal{E} denote a pairwise disjoint collection of edge-compressing disks incident to e . Isotope the disks in \mathcal{E} so that for each $E \in \mathcal{E}$, there is a distinct $t_E \in I$ such that $E \cap R_e \subset t_E \times I$. For each disk $E \in \mathcal{E}$, let γ_E^\pm denote the arcs of $H \cap R_e$ incident to $E \cap R_e$. We say E is a *leftmost* disk of \mathcal{E} if for each $E' \in \mathcal{E}$ that meets γ_E^+ or γ_E^- , $t_E < t_{E'}$. See Figure 4. Rightmost disks are defined similarly.




FIGURE 2. Defining the push-off E' of an edge-compressing disk E .





FIGURE 3. The second push-off E'' of E .

Note that it follows immediately from the above definitions that if E is a leftmost disk of a collection \mathcal{E} , and E' is the push-off of E to its left, then E' is disjoint from every disk in \mathcal{E} .


 FIGURE 4. Defining *leftmost* and *rightmost*.

As H has local index n , there is a homotopically non-trivial map ι from an $(n-1)$ -sphere S into $[\mathcal{CE}(H)]$. Triangulate S so that the map ι is simplicial.

Definition 4.6. Let e be an edge of Δ . An e -simplex of S is a simplex spanned by vertices that represent edge-compressing disks incident to e . Such a simplex is *full* if it is not a face of another e -simplex.

Lemma 4.7. S is homotopic to a sphere that contains no vertices representing edge-compressing disks incident to any short or medium edge.

Proof. Choose ι and the triangulation of S so that the sum ϵ of the total number of s -, s' -, m -, and m' -simplices is minimal. Our goal is to show that $\epsilon = 0$. If not, then S contains an e -simplex for some short or medium edge e . Thus, S contains a full e -simplex, σ . Let \mathcal{E} denote a collection of disks represented by the vertices of σ .

Note that it follows from the fact that e is a short or medium edge, and that there are no triangular components of ∂H , that there are no switches on e . Hence, if \mathcal{E} contains a disk incident to e then any rightmost disk E of \mathcal{E} on e has a push-off to the right. Such a push-off E' will be incident to a long edge and will be disjoint from every disk in \mathcal{E} .

Let Σ denote the union of the simplices in S that contain σ . Suppose $[X]$ is a vertex of Σ that is not in σ . If X meets the edge e , then the

join of σ with $[X]$ would be an e -simplex containing σ as a face. Since σ is full, no such simplex exists. We conclude X does not meet e . Furthermore, as $[X] \in \Sigma$, it is connected to every vertex of σ by an edge. Thus, the disk X is disjoint from every element of \mathcal{E} . It follows that X will be disjoint from the push-off E' , as this disk is constructed from E (which is disjoint from X) and a subset of $\partial\Delta$.

As E' is disjoint from every element of \mathcal{E} and every disk represented by a vertex of $\Sigma - \sigma$, we conclude the vertex $[E']$ is connected by an edge to every vertex of Σ in $[\mathcal{CE}(H)]$. Thus, we may replace Σ in S with the cone of $\partial\Sigma$ to the point $[E']$. Since σ is no longer an e -simplex of the resulting sphere, and no new simplices have been created spanned by edge-compressing disks incident to any other short or medium edge, we have lowered ϵ . \square

By Lemma 4.7, we may henceforth assume that S has no vertices representing edge-compressing disks incident to any short or medium edge. We now assume that in addition, S has been chosen so as to minimize the total number δ of l -simplices.

Lemma 4.8. *There is an l -simplex of S , and every full such simplex contains vertices representing edge-compressing disks that meet the points of $H \cap l$ that are closest to the endpoints of l .*

Proof. Let v^* denote one of the endpoints of l . Let Δ_{v^*} denote the union of the three 2-simplices of $\partial\Delta$ that meet v^* . Let s^* , m^* , and l^* denote the points of ∂H that are closest to v^* , on the short, medium, and long edges incident to v^* . Let E be the closure of the component of $\Delta_{v^*} - H$ that contains the vertex v^* . See Figure 5. Then pushing E slightly into Δ turns it into an edge-compressing disk for H incident to the edge l' opposite l . As E is isotopic into $\partial\Delta$, it is disjoint from every compressing disk for H , and every edge-compressing disk for H that does not meet s^* , m^* , or l^* . Thus, if S contains no vertices representing edge-compressing disks that meet s^* , m^* , and l^* , then $\iota(S)$ is homotopic to the point $[E]$, a contradiction. However, by assumption, there are no edge-compressing disks represented by vertices of S that are incident to either s or m . We conclude there is a vertex of S that represents an edge-compressing disk incident to l^* .

The above argument implies there is a vertex of S representing a disk that meets the edge l . Hence, there is an l -simplex, and thus a full l -simplex, σ . Let Σ denote the union of the simplices of S that contain σ . Then no vertex of $\Sigma - \sigma$ represents a disk incident to l , and hence all such vertices are connected by an edge to $[E]$ in $[\mathcal{CE}(H)]$. Similarly, if no vertex of σ represents a disk that meets the point l^* ,




FIGURE 5. The three 2-simplices of $\partial\Delta$ that meet v^* , and the edge-compressing disk E incident to l' .

then all such vertices are also connected by an edge to $[E]$ in $[\mathcal{CE}(H)]$. We conclude that in this case we will be able to lower δ by replacing Σ with the cone on $\partial\Sigma$ to the point $[E]$. This contradiction establishes that any full l -simplex of S contains an edge-compressing disk incident to l^* . \square

Theorem 4.9. *If H is an index n topologically minimal surface in a tetrahedron where every component of ∂H has length at least 8, then $|\partial H| \leq 4(n+1)$.*

Note that when H is simply connected, then in [Bacb] we have already established that $|\partial H|$ is precisely $4(n+1)$.

Proof. By Lemma 4.8, there is a full l -simplex σ of S . Let \mathcal{E} denote a collection of disks representing the vertices of σ . We claim that each point of l between two points of $H \cap l$ meets some disk in \mathcal{E} . Suppose not. Let p be a point of l between points of $H \cap l$ that does not meet any disk in \mathcal{E} . By Lemma 4.8, there are disks in \mathcal{E} that meet l on both sides of p . Thus, there is both a leftmost and a rightmost disk of \mathcal{E} on both sides of p . As there is only one left switch on l , there must then be a leftmost disk of \mathcal{E} that does not meet the left switch. Such a disk E will have a second push-off to a disk E' that meets the edge l' , where E' is disjoint from every disk in \mathcal{E} .

The remainder of the proof is similar to that of Lemma 4.8. Let Σ denote the union of the simplices of S that contain σ . Then by construction, every vertex of Σ is connected by an edge to $[E]$ in $[\mathcal{CE}(H)]$. We may therefore lower $c(S; l)$ by replacing Σ with the cone on $\partial\Sigma$ to the point $[E]$.

Since n is the index of H , the dimension of S is $n - 1$. Thus, every simplex of S is spanned by at most n vertices. Hence, $|\mathcal{E}| \leq n$. The points of $H \cap l$ separate l into $|H \cap l| + 1$ intervals. By the above argument, all but the first and last such intervals meets a disk in \mathcal{E} . Thus, $|H \cap l| + 1 - 2 \leq n$, or $|H \cap l| \leq n + 1$. As l is a long edge, $|\partial H| = 4|H \cap l|$. Putting this together then gives us $|\partial H| \leq 4(n + 1)$, as desired. \square

The following lemma takes us one more step closer to establishing the third conclusion of Theorem 1.1.

Lemma 4.10. *If H is a surface in a tetrahedron with well-defined local index, and every component of ∂H has length at least 8, then H is connected.*

Proof. Let S be as above. According to the proof of Theorem 4.9, each point p of l between two points of $H \cap l$ meets some edge-compressing disk E represented by a vertex of S . Thus, $E \cap H$ is a path in H connecting the points of $H \cap l$ on either side of p . It follows that these points are on the same component of H . As this is true for each consecutive pair of points of $H \cap l$, we conclude that all of these points are on the same component of H . The result thus follows. \square

5. COMPRESSING AWAY TRIANGULAR BOUNDARY COMPONENTS

In this section we assume H is normally connected, and thus has no components that are normal triangles or quadrilaterals. Let v be a vertex of Δ , and suppose ∂H has a component of length 3 which separates v from the other boundary components of H . This length 3 component bounds a subdisk C_0 of $\partial\Delta$ than contains v . As H has no components that are normal triangles, C_0 is isotopic to a compressing disk for H .

Lemma 5.1. *Every essential sphere S in $[\mathcal{CE}(H)]$ has a vertex that represents an edge-compressing disk that meets C_0 .*

Proof. As C_0 is isotopic into $\partial\Delta$ it is disjoint from every other compressing disk for H . If it is also disjoint from every edge-compressing disk represented by a vertex of S , then S is homotopic to the point $[C_0]$ in $[\mathcal{CE}(H)]$. \square




FIGURE 6. The rectangle $Q \subset \partial\Delta^1$ when E_1 and E_2 are adjacent neighbors. In the figure on the left, E_1 and E_2 are incident to the same edge of Δ . In the figure on the right, they are incident to different edges.

Let Δ_v denote the union of the 2-simplices of Δ that contain v . Let Δ^1 denote Δ with a neighborhood of the 1-skeleton removed, and $H^1 = H \cap \Delta^1$. Let S be an essential sphere in $[\mathcal{CE}(H)]$. Let $[E_0]$ be a vertex of S , where E_0 is an edge-compressing disk that meets C_0 (as guaranteed to exist by Lemma 5.1). Let \mathcal{C}_0 denote the set which contains the single element C_0 . Let \mathcal{E}_0 be a set of edge-compressing disks that meets C_0 , includes E_0 , and represents the vertices of a simplex of S .

Now suppose $E_1, E_2 \in \mathcal{E}_0$. Then E_1 and E_2 each meet ∂C_0 in a point. If the subarc of ∂C_0 connecting these points does not meet any other disk in \mathcal{E}_0 , then we say E_1 and E_2 are *consecutive*. If the subarc of ∂C_0 connecting consecutive disks only meets the edges of Δ that either E_1 or E_2 meet, then we say E_1 and E_2 are *neighbors*. Neighboring disks E_1 and E_2 each meet some other points of ∂H^1 . If these two points are connected by a subarc of ∂H^1 that lies in Δ_v , then we say E_1 and E_2 are *adjacent*. Otherwise, E_1 and E_2 are *non-adjacent*.

When E_1 and E_2 are adjacent neighbors, then the arcs $E_1 \cap \partial\Delta^1$ and $E_2 \cap \partial\Delta^1$, together with two subarcs of ∂H^1 , bound a rectangle Q in $\partial\Delta^1$, as in Figure 6. If the disk $C = E_1 \cup Q \cup E_2$ is a compressing disk for H , then we say C is *defined* by E_1 and E_2 .

When E_1 and E_2 are non-adjacent neighbors, then the arcs $E_1 \cap \partial\Delta^1$ and $E_2 \cap \partial\Delta^1$, together with three subarcs of ∂H^1 and a subarc of $\partial\Delta_v$, bounds a hexagonal region Q of $\partial\Delta^1$, as in Figure 7. In this case, the disk $C = E_1 \cup Q \cup E_2$ is an edge-compressing disk for H , that we say is *defined* by E_1 and E_2 .





FIGURE 7. The hexagon $Q \subset \partial\Delta^1$ when E_1 and E_2 are non-adjacent neighbors.

We now let \mathcal{C}_1 denote the set of all compressing disks for H defined by pairs of adjacent neighbors in \mathcal{E}_0 , together with the set of all edge-compressing disks defined by non-adjacent neighbors in \mathcal{E}_0 . Let \mathcal{E}_1 be a set of edge-compressing disks such that

- for each $E \in \mathcal{E}_1$ there is a $C \in \mathcal{C}_1$ such that $E \cap C \neq \emptyset$, and
- the disks of $\mathcal{E}_0 \cup \mathcal{E}_1$ are represented by the vertices of a simplex in S .

For example, Figure 8 depicts a surface H and a set \mathcal{E}_0 with three disks. Three pairs of these disks are consecutive, two of those pairs are neighbors, and one of those neighbors defines the compressing disk for the lower tube in the figure, which would then be an element of \mathcal{C}_1 . Also depicted in the figure is a single edge-compressing disk that meets this compressing disk, which is hence an element of \mathcal{E}_1 . Note that all the disks in $\mathcal{E}_0 \cup \mathcal{E}_1$ are disjoint, and so are represented by the vertices of a simplex in $[\mathcal{CE}(H)]$.

We now proceed in the same fashion. Two disks of \mathcal{E}_1 are consecutive if they meet the same disk in \mathcal{C}_1 , and there are no other disks of \mathcal{E}_1 between them. The terms *neighbor*, *adjacent* and *non-adjacent* are defined as before, as are compressing disks defined by adjacent neighbors, and edge-compressing disks defined by non-adjacent neighbors. In this way we construct sequences of sets $\{\mathcal{C}_0, \mathcal{C}_1, \mathcal{C}_2, \dots\}$ and $\{\mathcal{E}_0, \mathcal{E}_1, \mathcal{E}_2, \dots\}$, such that for each $i \geq 1$,


 FIGURE 8. Three disks in \mathcal{E}_0 and one in \mathcal{E}_1 .

- the set \mathcal{C}_i consists of all compressing disks defined by adjacent neighbors in \mathcal{E}_{i-1} , and edge-compressing disks defined by non-adjacent neighbors in \mathcal{E}_{i-1} ,
- the set \mathcal{E}_i consists of edge-compressing disks such that for each $E \in \mathcal{E}_i$ there is a $C \in \mathcal{C}_i$ where $E \cap C \neq \emptyset$, and
- the disks of $\bigcup_{j=0}^i \mathcal{E}_j$ are represented by the vertices of a simplex in S .

Definition 5.2. A simplex $\sigma \subset S$ whose vertices represent the disks $\bigcup \mathcal{E}_i$ as defined above is called a *sunburst* of S .

Definition 5.3. Suppose σ is a sunburst and $\{\mathcal{E}_i\}_{i=0}^n$ is the sequence of sets of disks represented by its vertices. Then the *sunburst complexity* of σ is the $(n+1)$ -tuple

$$c(\sigma) = (|\mathcal{E}_0|, |\mathcal{E}_1|, \dots, |\mathcal{E}_n|).$$

Two such $(n+1)$ -tuples are compared lexicographically.

Definition 5.4. When σ is a sunburst that is not a face of any other sunburst, then we say it is *full*.

Definition 5.5. The *sunburst complexity* of S is defined to be the set of $(n+1)$ -tuples

$$\{c(\sigma) | \sigma \text{ is a full sunburst of } S\}.$$

We order this set in non-increasing order, and compare two such sets lexicographically. We say an essential sphere S is *sunburst minimal* if

there is no sphere S' homotopic to S that has lower sunburst complexity.

Lemma 5.6. *Suppose S is sunburst minimal, σ is a full sunburst of S , $\{\mathcal{E}_i\}$ is the sequence of sets of disks represented by its vertices, and for each i , \mathcal{C}_i is the set of compressing and edge-compressing disks defined by neighboring disks in \mathcal{E}_i . Then for each $i \geq 1$ and each disk $C \in \mathcal{C}_i$, there is a disk $E \in \mathcal{E}_i$ that meets C .*

Proof. Let Σ be the union of simplices in S that contain σ . A vertex $[F]$ in $\Sigma - \sigma$ represents a disk F that is disjoint from every element of $\bigcup \mathcal{E}_i$. Suppose there is an i such that \mathcal{C}_i contains a disk C that is disjoint from F . Then adding F to \mathcal{E}_i defines a new sunburst which is the join of σ and $[F]$, contradicting our assumption that σ was full. We conclude that each disk $C \in \mathcal{C}_i$ is disjoint from every disk represented by a vertex of $\Sigma - \sigma$.

Now suppose C is a disk in \mathcal{C}_i that does not meet any disk in \mathcal{E}_i . Such a disk C will then be disjoint from every element of $\bigcup \mathcal{E}_i$, and thus C is disjoint from every disk represented by a vertex of Σ . We may thus obtain a sphere homotopic to S by replacing Σ with the cone of $\partial\Sigma$ to the point $[C]$, $\partial\Sigma * [C]$. Since C is a compressing disk, any full sunburst in $\partial\Sigma * [C]$ will have fewer edge-compressing disks, and thus a smaller sunburst complexity. \square

Lemma 5.7. *Suppose S is sunburst minimal, σ is a full sunburst of S , and $\{\mathcal{E}_i\}$ is the sequence of sets of disks represented by its vertices. Then for each i , all neighbors in \mathcal{E}_i are adjacent.*

Proof. Suppose E_1 and E_2 are non-adjacent neighbors in \mathcal{E}_i , and let C denote the edge-compressing disk in \mathcal{C}_{i+1} that they define. By Lemma 5.6, C meets some disk $E \in \mathcal{E}_{i+1}$. Since E_1 and E_2 are non-adjacent, they must be incident to different edges of Δ . The disk E will then be incident to the same edge as E_1 , or the same edge as E_2 . Assume the former, and let e denote the edge of Δ that both E_1 and E meet.

Comparing Figures 1 and 7, we see that E_1 meets R_e (the frontier of $N(e)$) in either a right or left switch. Note that since E is disjoint from E_1 and E_2 , then it must meet the disk Q , where $C = E_1 \cup Q \cup E_2$ and Q is the shaded region of Δ^1 depicted in Figure 7. The three simplices of $\partial\Delta^1$ that are incident to the vertex v are depicted in Figure 9, with Q depicted again as a shaded region. From the figure, we see that the disk E is to the right of E_1 on e . Hence, the simplex σ will have a rightmost disk E' on e which has a push-off E^* to its right, which is an edge-compressing disk incident to an edge that does not meet the




FIGURE 9. The disk E is further right than the disk E_1 , which meets the right switch.

vertex v . By construction then, E^* will be disjoint from every disk represented by a vertex of σ .

Let Σ be the union of simplices in S that contain σ . Let $[F]$ be a vertex of $\Sigma - \sigma$. As E^* is a push-off, it consists of the disk E' , together with a rectangular subregion of $\partial\Delta^1$, as depicted in Figure 2. Hence, if F meets E^* it must also meet E' . However, $[F]$ is by definition connected by an edge to each vertex of σ , and thus F is disjoint from every disk represented by a vertex of σ . In particular, F is disjoint from E' . We conclude F must also be disjoint from E^* .

We conclude E^* is disjoint from every disk represented by a vertex of Σ . We may thus obtain a sphere homotopic to S by replacing Σ with the cone of $\partial\Sigma$ to the point $[E^*]$. Every full sunburst simplex of such a cone is spanned a proper subset of the vertices of σ . Hence, any resulting full sunburst will have lower sunburst complexity. \square

Theorem 5.8. *Suppose S is sunburst minimal, σ is a full sunburst of S , $\{\mathcal{E}_i\}$ is the sequence of sets of disks represented by its vertices, and for each i , \mathcal{C}_i is the set of compressing disks defined by neighboring disks in \mathcal{E}_i . Then the local index of $H/\bigcup \mathcal{C}_i$ is at most $n - |\bigcup \mathcal{E}_i|$.*

Proof. Let Σ be the union of simplices in S that contain σ . We claim $\Sigma - \sigma$ is an essential sphere in $[\mathcal{CE}(H/\bigcup \mathcal{C}_i)]$. To establish this, there are a number of things we will have to prove.

Claim 5.9. *Each vertex in $\Sigma - \sigma$ represents a disk in $\mathcal{CE}(H/\bigcup \mathcal{C}_i)$.*

Proof. Suppose $[D]$ is such a vertex, representing a disk, D . We first deal with the case that D is an edge-compressing disk. If such a disk is disjoint from every disk in \mathcal{C}_i , for each i , then it will be an edge-compressing disk for $H/\bigcup \mathcal{C}_i$. However, if D meets some disk $C \in \mathcal{C}_i$, for some i , then D would have belonged to the set \mathcal{E}_{i+1} , and thus $[D]$ would have been a vertex of σ (since σ is assumed to be full).

We now deal with the case that D is a compressing disk. We again claim that it must be disjoint from every disk in \mathcal{C}_i , for each i . Suppose $C \in \mathcal{C}_i$ is a disk such that $D \cap C \neq \emptyset$. Recall that C is either the disk C_0 , or by Lemma 5.7, $C = E_1 \cup Q \cup E_2$ for some rectangle $Q \subset \Delta^1$ and adjacent disks $E_1, E_2 \in \mathcal{E}_{i-1}$. In the former case, the disk C is isotopic into $\partial\Delta^1$, and is therefore disjoint from every compressing disk such as the disk D . In the latter case, we know D is disjoint from E_1 and E_2 , since $[D]$ is a vertex in the link of σ , and $[E_1]$ and $[E_2]$ are vertices of σ . Furthermore, D is disjoint from Q , as $Q \subset \partial\Delta^1$. Therefore D is disjoint from $C = E_1 \cup Q \cup E_2$.

We have established that D is disjoint from every disk in \mathcal{C}_i , for all i . To show D is a compressing disk for $H/\bigcup \mathcal{C}_i$, we must now show that ∂D is essential on $H/\bigcup \mathcal{C}_i$. If not, then ∂D bounds a subdisk D' of $H/\bigcup \mathcal{C}_i$ that contains scars of a disk $C \in \mathcal{C}_i$, for some i . (See Definition 2.4.) By Lemma 5.6, there is an edge-compressing disk $E \in \mathcal{E}_i$ for H that meets C . Thus, $E \cap \partial D \neq \emptyset$, and hence $E \cap D \neq \emptyset$. We have reached a contradiction, as D is disjoint from every disk in \mathcal{E}_i , for all i . \square

The previous claim implies that there is a natural map ϕ from $S' = \Sigma - \sigma$ into $[\mathcal{CE}(H/\bigcup \mathcal{C}_i)]$. We must now show that this map is homotopically non-trivial. If not, then we can extend ϕ to a map from a ball B bounded by S' into $[\mathcal{CE}(H/\bigcup \mathcal{C}_i)]$. We assume B has been triangulated so that this map is simplicial.

We now construct a map $\eta : B \rightarrow [\mathcal{CE}(H)]$ such that

- (1) $\eta \circ \phi$ fixes S' , and
- (2) each disk represented by a vertex of $\eta(B)$ is disjoint from every disk represented by a vertex of σ .

To prove the existence of η we must show that for each disk $D \in \mathcal{CE}(H/\bigcup \mathcal{C}_i)$ there is a disk $\tilde{D} \in \mathcal{CE}(H)$, disjoint from every disk represented by a vertex of σ , that becomes D after compressing. It will follow from this construction that if D and D' are disjoint disks in $\mathcal{CE}(H/\bigcup \mathcal{C}_i)$ then \tilde{D} and \tilde{D}' can be chosen to be disjoint. Thus, the map η that send D to \tilde{D} will map B simplicially into $[\mathcal{CE}(H)]$.

The map η will be easily defined once we better understand the relationship between the surfaces H and $H/\bigcup \mathcal{C}_i$. When H is the surface




FIGURE 10. The surface $H/\bigcup \mathcal{C}_i$, when H is the surface of Figure 8.

depicted in Figure 8, the surface $H/\bigcup \mathcal{C}_i$ is depicted on the left in Figure 10. The small circles in this figure bound disks that are the scars of the disks in $\mathcal{C}_0 \cup \mathcal{C}_1$. The arcs that meet these disks represent the places where the disks of $\mathcal{E}_0 \cup \mathcal{E}_1$ meet H . On the righthand side of this figure we depict the following: The circle at the top depicts the upper scar of \mathcal{C}_0 . The three arcs incident to it represent places where the three disks of \mathcal{E}_0 meet H . Two pairs of these arcs represent neighboring disks in \mathcal{E}_0 , and are thus connected by subarcs of ∂H . Just below this we depict the other scar of \mathcal{C}_0 . Again there are three arcs incident to this scar, also representing places where the three disks of \mathcal{E}_0 meet H . As in the upper figure, for each pair of neighboring disks in \mathcal{E}_0 there is a subarc of ∂H depicted joining them. Furthermore, one such pair defines a compressing disk \mathcal{C}_1 in \mathcal{C}_1 . A scar of \mathcal{C}_1 is depicted by the smaller circle of the middle figure, together with an arc that represents where a disk of \mathcal{E}_1 meet H . Finally, the lowest figure depicts the other scar of \mathcal{C}_1 , together with the other arc where a disk of \mathcal{E}_1 meet H .

We call the entire righthand figure a *sunburst diagram*. Formally, a sunburst diagram is a depiction of each scar of a disk of $\bigcup \mathcal{C}_i$, together with arcs representing all of the places where disks of $\bigcup \mathcal{E}_i$ meet H , along with subarcs of ∂H connecting arcs representing neighboring disks of $\bigcup \mathcal{E}_i$. Note that between any pair of neighboring disks such a subarc of ∂H must exist, because otherwise such a pair would be non-adjacent. By Lemma 5.7, there are no such pairs.




FIGURE 11. A more complicated sunburst diagram.

Sunburst diagrams can be far more complicated than the one depicted in Figure 10. Knowing what features such a diagram must necessarily possess will be helpful in understanding the rest of the proof. As an illustrative example, consider Figure 11. Such a diagram comes from a situation in which there are four disks in \mathcal{E}_0 . At least one of these disks meets each edge of Δ that is incident to v , and hence every pair of consecutive disks are neighbors. By Lemma 5.7, such neighbors must be adjacent, so there is an arc of ∂H connecting them. This is why the boundaries of the first and second figure in the diagram are complete circles. In the figure, two neighboring pairs of \mathcal{E}_0 define two compressing disks in \mathcal{C}_1 . Each of the disks in \mathcal{C}_1 meets three disks in \mathcal{E}_1 . Three pairs of the disks in \mathcal{E}_1 define three compressing disks in \mathcal{C}_2 . Two of these meet a single disk in \mathcal{E}_2 , and one meets two disks in \mathcal{E}_2 . Finally, the only two neighboring disks in \mathcal{E}_2 are parallel on H , and therefore do not define any further compressing disks.

Claim 5.10. *The subset of ∂H represented in a component of a sunburst diagram is connected, and is thus either an arc or a circle. In the former case, the remnants of $\bigcup \mathcal{E}_i$ depicted in this component are incident to at most two edges of Δ . In the latter case, the depicted circle of ∂H bounds a component of $H/\bigcup \mathcal{C}_i$ that is a normal triangle.*

Proof. Each arc in a sunburst diagram represents a disk of \mathcal{E}_i that is incident to one of the three edges of Δ that meet v . If all three edges of Δ are met, then each pair of consecutive disks of \mathcal{E}_i represented in the diagram are neighbors. By Lemma 5.7, such a pair is adjacent, and thus they are connected by an arc of ∂H . As there are three edges of Δ incident to v , these arcs piece together to form a length three normal loop on $\partial \Delta$. After compressing along the disks of $\bigcup \mathcal{C}_i$, this loop will bound a normal disk.




FIGURE 12. On the left is a component of a sunburst diagram that meets a compressing disk D . Since the depicted subset of ∂H is a connected arc, ∂D can be isotoped to go around this component of the sunburst diagram, as in the figure on the right.

If at most two of the edges of Δ are met by disks of \mathcal{E}_i represented in a component of a sunburst diagram, then exactly one pair of consecutive disks of \mathcal{E}_i are not neighbors. Hence, all of the other pairs are joined by arcs of ∂H to form a connected arc. \square

We now complete the proof of Theorem 5.8. Suppose first D is a compressing or edge-compressing disk for $H/\bigcup \mathcal{C}_i$. As each of the scars of the disks in $\bigcup \mathcal{C}_i$ is a disk, ∂D can be isotoped to be disjoint from these. If ∂D meets any of the remnants of the disks of $\bigcup \mathcal{E}_i$, then ∂D will be depicted in some component of a sunburst diagram. The subset of ∂H depicted in this component can not be a circle, since Claim 5.10 would then imply that this component represents a normal triangle of $H/\bigcup \mathcal{C}_i$. However, no compressing or edge-compressing disk for $H/\bigcup \mathcal{C}_i$ will meet a normal triangle piece of this surface.

Claim 5.10 now implies that the subset of ∂H depicted in any component of the sunburst diagram that meets ∂D is an arc. Hence, this component looks something like the one depicted in Figure 12. The salient feature of this figure is that there is a (larger) circle representing a scar of some disk C of \mathcal{C}_i (for some i), arcs of \mathcal{E}_i connecting this circle to a connected subarc of ∂H , and that the entire figure is homeomorphic to a disk. Thus, any intersection of ∂D with this component can be isotoped off. When D is a compressing disk, such an isotopy is




FIGURE 13. A component of a sunburst diagram that meets an edge-compressing disk D . The disk E_1 is incident to the same edge of Δ as D .

depicted in Figure 12. When D is an edge-compressing disk some additional care must be taken. The depicted disks of \mathcal{E}_i are $\{E_j\}_{j=1}^m$, where for each j , the disks E_j and E_{j+1} are adjacent neighbors. These disks are all incident to at most two edges of Δ . Hence, we may assume that for some $k \leq m$, the disks $\{E_j\}_{j=1}^k$ are all incident to the same edge of Δ . The disk D is then either incident to the same edge as these disks, or $k < m$ and D is incident to the same edge as the disks $\{E_j\}_{j=k+1}^m$. Assume the former. Then ∂D can be isotoped to be disjoint from every disk in $\{E_j\}_{j=1}^m$ via an isotopy as depicted in Figure 13.

Once ∂D has been isotoped to be disjoint from every component of a sunburst diagram, it follows from a standard innermost disk argument that the disk D can be made disjoint from every disk of $\bigcup \mathcal{C}_i$ and $\bigcup \mathcal{E}_i$. Thus, D is now the desired compressing or edge-compressing disk \tilde{D} for $H/\bigcup \mathcal{C}_i$, and is disjoint from every disk represented by a vertex of σ .

The proof is now complete by observing that the existence of the ball B and map η implies that we can modify S by replacing Σ with the join of $\eta(B)$ and $\partial\sigma$, as in Figure 14. This will lower the sunburst complexity of S , providing the desired contradiction. \square

The following lemma completes the proof of the third and fourth conclusion of Theorem 1.1.

Lemma 5.11. *Suppose H is a normally connected surface in a tetrahedron with local index n . Then H is connected, and $|\partial H| \leq 4(n+1)$.*




FIGURE 14. In the left is the simplex σ , and the simplices Σ containing σ . On the right is the join of $\eta(B)$ with $\partial\sigma$.

Proof. First, suppose every component of ∂H has length three or four. By Lemma 3.6, H is then connected. By the second conclusion of Theorem 1.1, $1 - \chi(H) \leq n$. It follows that H has at most $n + 1$ boundary components. By assumption, each boundary component has length at most four, so we have $|\partial H| \leq 4(n + 1)$.

When every component of ∂H has length at least eight, then H is connected by Lemma 4.10 and $|\partial H| \leq 4(n + 1)$ by Theorem 4.9.

We are left with the case that ∂H has some components of length at least 8, and some components of length 3. It follows from Lemma 3.5 that all parallel length 3 components are on the same component of H , and each such parallel family is on the same component of H as at least one longer component of ∂H . We will show that all of these longer components of ∂H are on the same component of H , and thus conclude that H must be connected.

Let $\{\mathcal{E}_i\}$ and $\{\mathcal{C}_i\}$ denote the collections of disks of Theorem 5.8. By Lemma 5.6, $|\bigcup \mathcal{E}_i| \geq |\bigcup \mathcal{C}_i|$. By Theorem 5.8, the local index of $H/\bigcup \mathcal{C}_i$ is at most $n - |\bigcup \mathcal{E}_i|$, and thus at most $n - |\bigcup \mathcal{C}_i|$. By construction, the set \mathcal{C}_0 contains a compressing disk C_0 such that H/C_0 contains a component that is a normal triangle. Thus, $H/\bigcup \mathcal{C}_i$ contains at least one normal triangle component. Furthermore, the total number t of normal triangle components of $H/\bigcup \mathcal{C}_i$ is at most $|\bigcup \mathcal{C}_i|$. Thus, the local index of $H/\bigcup \mathcal{C}_i$ is at most $n - t$.

We now continue to apply Theorem 5.8 to obtain a sequence of surfaces. Each such surface is obtained from the previous one by some compressions, followed by discarding any resulting normal triangle components. This process concludes with a surface H' with only boundary curves of length at least eight. By the preceding paragraph, the local

index n' of H' will be at most n minus the total number of normal triangles lost, which is precisely the number m of length three components of ∂H . That is, $n' \leq n - m$. As above, it follows from Lemma 4.10 that H' is connected, and from Theorem 4.9 that $|\partial H'| \leq 4(n' + 1)$. Thus, H was connected, and

$$\begin{aligned} |\partial H| &= |\partial H'| + 3m \\ &\leq 4(n' + 1) + 3m \\ &\leq 4(n - m + 1) + 3m \\ &\leq 4(n + 1). \end{aligned}$$

□

6. QUESTIONS AND CONJECTURES

Theorem 1.1 narrows down the possible pictures for a surface with a prescribed local index in a tetrahedron to a finite list. However, as noted in the introduction, this list will almost certainly contain surfaces that do not have the desired local index. Resolving the following questions would help give an exact classification.

For all of the following questions, H will denote a connected surface with local index n in a tetrahedron Δ .

Question 6.1. *Can ∂H have more than one “long” component?*

Here, a “long component” is one that meets the 1-skeleton at least eight times. In particular, it does not seem likely that two parallel n -gons, connected by a single unknotted tube, will be locally topologically minimal.

Question 6.2. *Does every long component of ∂H have to be the boundary of a helicoid?*

By [Bach], if H is a disk, then it is a helicoid. This rules out, for example, certain loops of length 20 from being the boundary of a locally topologically minimal disk. Can such loops arise as boundary components of locally topologically minimal surfaces that are not disks?

A more careful analysis of the proof of Theorem 4.9 reveals that when the boundary of H is not helical, then the bound given in that theorem can be improved to $|\partial H| \leq 2n + 4$. As the shortest possible non-helical loop on $\partial\Delta$ has length 20, we conclude that the lowest possible local index of a surface with non-helical boundary is 8.

Question 6.3. *If ∂H contains a triangular component then does the surface obtained by compressing H along an outermost one reduce the index of H by exactly one?*

Theorem 5.8 implies that in this case there will be a collection of compressing disks that will pinch off this triangle, and lower the index by some amount. Can one improve the proof so that this collection always has exactly one element?

Question 6.4. *Can one determine the local index of H , just from knowing its Euler characteristic and total boundary length?*

We would like to think of local index as a commodity. Theorem 1.1 says that you can pay for this commodity with either negative Euler characteristic or boundary length. This question asks whether or not there is an exchange rate between these two currencies.

REFERENCES

- [Baca] D. Bachman. Normalizing Topologically Minimal Surfaces I: Global to Local Index. Preprint.
- [Bacb] D. Bachman. Normalizing Topologically Minimal Surfaces II: Disks. Preprint.
- [Bac04] David Bachman. A note on Kneser-Haken finiteness. *Proc. Amer. Math. Soc.*, 132(3):899–902 (electronic), 2004.
- [Bac10] D. Bachman. Topological Index Theory for surfaces in 3-manifolds. *Geometry & Topology*, 14(1):585–609, 2010.
- [Bac13] David Bachman. Stabilizing and destabilizing Heegaard splittings of sufficiently complicated 3-manifolds. *Math. Ann.*, 355(2):697–728, 2013.
- [BDTS] D. Bachman, R. Derby-Talbot, and E. Sedgwick. Heegaard structure respects complicated JSJ decompositions. Available at <http://arxiv.org/abs/0911.5078>.
- [CG87] A. J. Casson and C. McA. Gordon. Reducing Heegaard splittings. *Topology and its Applications*, 27:275–283, 1987.
- [CM04a] Tobias H. Colding and William P. Minicozzi, II. The space of embedded minimal surfaces of fixed genus in a 3-manifold. I. Estimates off the axis for disks. *Ann. of Math. (2)*, 160(1):27–68, 2004.
- [CM04b] Tobias H. Colding and William P. Minicozzi, II. The space of embedded minimal surfaces of fixed genus in a 3-manifold. II. Multi-valued graphs in disks. *Ann. of Math. (2)*, 160(1):69–92, 2004.
- [CM04c] Tobias H. Colding and William P. Minicozzi, II. The space of embedded minimal surfaces of fixed genus in a 3-manifold. III. Planar domains. *Ann. of Math. (2)*, 160(2):523–572, 2004.
- [CM04d] Tobias H. Colding and William P. Minicozzi, II. The space of embedded minimal surfaces of fixed genus in a 3-manifold. IV. Locally simply connected. *Ann. of Math. (2)*, 160(2):573–615, 2004.
- [Hak61] W. Haken. Theorie der Normalfächen. *Acta Math.*, 105:245–375, 1961.
- [Kne29] H. Kneser. Geschlossene Flächen in dreidimensionalen Mannigfaltigkeiten. *Jahresbericht der Dent. Math. Verein*, 28:248–260, 1929.
- [Li07] Tao Li. Heegaard surfaces and measured laminations. I. The Waldhausen conjecture. *Invent. Math.*, 167(1):135–177, 2007.

- [Rub95a] J. H. Rubinstein. An algorithm to recognise the 3-sphere. In *Proceedings of the International Congress of Mathematicians, Vol. 1, 2 (Zürich, 1994)*, pages 601–611. Birkhäuser, Basel, 1995.
- [Rub95b] J. H. Rubinstein. Polyhedral minimal surfaces, Heegaard splittings and decision problems for 3-dimensional manifolds. In *Proceedings of the Georgia Topology Conference*, pages 1–20, 1995.
- [Sto00] M. Stocking. Almost normal surfaces in 3-manifolds. *Trans. Amer. Math. Soc.*, 352:171–207, 2000.
- [Tho94] A. Thompson. Thin position and the recognition problem for the 3-sphere. *Math. Research Letters*, 1:613–630, 1994.